arXiv:1009.3078v1 [cs.CV] 16 Sep 2010

Asymmetric Totally-corrective Boosting for Real-time Object Detection

Peng Wang[1*], Chunhua Shen[2,3], Nick Barnes[2,3], Hong Zheng[1], Zhang Ren[1]

[1] Beihang University, Beijing 100191, China
[2] NICTA[†], Canberra Research Laboratory, Canberra, ACT 2601, Australia
[3] Australian National University, Canberra, ACT 0200, Australia

Abstract. Real-time object detection is one of the core problems in computer vision. The cascade boosting framework proposed by Viola and Jones has become the standard for this problem. In this framework, the learning goal for each node is asymmetric, which is required to achieve a high detection rate and a moderate false positive rate. We develop new boosting algorithms to address this asymmetric learning problem. We show that our methods explicitly optimize asymmetric loss objectives in a totally corrective fashion. The methods are totally corrective in the sense that the coefficients of all selected weak classifiers are updated at each iteration. In contract, conventional boosting like AdaBoost is stage-wise in that only the current weak classifier's coefficient is updated. At the heart of the totally corrective boosting is the column generation technique. Experiments on face detection show that our methods outperform the state-of-the-art asymmetric boosting methods.

1 Introduction

Due to its important applications in video surveillance, interactive human-machine interface *etc*, real-time object detection has attracted extensive research recently [1–6]. Although it was introduced a decade ago, the boosted cascade classifier framework of Viola and Jones [2] is still considered as the most promising approach for object detection, and this framework is the basis which many papers have extended.

One difficulty in object detection is the problem is highly asymmetric. A common method to detect objects in an image is to exhaustively search all sub-windows at all possible scales and positions in the image, and use a trained model to detect target objects. Typically, there are only a few targets in millions of searched sub-windows. The cascade classifier framework partially solves the asymmetry problem by splitting the detection process into several nodes. Only those sub-windows passing through all nodes are classified as true targets. At each node, we want to train a classifier with a very high detection rate (*e.g.*, 99.5%) and a moderate false positive rate (*e.g.*, around

[*] Work was done when P. W. was visiting NICTA Canberra Research Laboratory and Australian National University.

[†] NICTA is funded by the Australian Government's Department of Communications, Information Technology, and the Arts and the Australian Research Council through *Backing Australia's Ability* initiative and the ICT Research Center of Excellence programs.

50%). The learning goal of each node should be asymmetric in order to achieve optimal detection performance. A drawback of standard boosting like AdaBoost in the context of the cascade framework is it is designed to minimize the overall false rate. The losses are equal for misclassifying a positive example and a negative example, which makes it not be able to build an optimal classifier for the asymmetric learning goal.

Many subsequent works attempt to improve the performance of object detectors by introducing asymmetric loss functions to boosting algorithms. Viola and Jones proposed asymmetric AdaBoost [3], which applies an asymmetric multiplier to one of the classes. However, this asymmetry is absorbed immediately by the first weak classifier because AdaBoost's optimization strategy is greedy. In practice, they manually apply the n-th root of the multiplier on each iteration to keep the asymmetric effect throughout the entire training process. Here n is the number of weak classifiers. This heuristic cannot guarantee the solution to be optimal and the number of weak classifiers need to be specified before training. AdaCost presented by Fan *et al.* [7] adds a cost adjustment function on the weight updating strategy of AdaBoost. They also pointed out that the weight updating rule should consider the cost not only on the initial weights but also at each iteration. Li and Zhang [8] proposed FloatBoost to reduce the redundancy of greedy search by incorporating floating search with AdaBoost. In FloatBoost, the poor weak classifiers are deleted when adding the new weak classifier. Xiao *et al.* [9] improved the backtrack technique in [8] and exploited the historical information of preceding nodes into successive node learning. Hou *et al.* [10] used varying asymmetric factors for training different weak classifiers. However, because the asymmetric factor changes during training, the loss function remains unclear. Pham *et al.* [11] presented a method which trains the asymmetric AdaBoost [3] classifiers under a new cascade structure, namely multi-exit cascade. Like soft cascade [12], boosting chain [9] and dynamic cascade [13], multi-exit cascade is a cascade structure which takes the historical information into consideration. In multi-exit cascade, the n-th node "inherits" weak classifiers selected at the preceding $n-1$ nodes. Wu *et al.* [14] stated that feature selection and ensemble classifier learning can be decoupled. They designed a linear asymmetric classifier (LAC) to adjust the linear coefficients of the selected weak classifiers. Kullback-Leibler Boosting [15] iteratively learns robust linear features by maximizing the Kullback-Leibler divergence.

Much of the previous work is based on AdaBoost and achieves the asymmetric learning goal by *heuristic* weights manipulations or post-processing techniques. It is not trivial to assess how these heuristics affect the original loss function of AdaBoost. In this work, we construct new boosting algorithms *directly* from asymmetric losses. The optimization process is implemented by column generation. Experiments on toy data and real data show that our algorithms indeed achieve the asymmetric learning goal without any heuristic manipulation, and can outperform previous methods.

Therefore, the main contributions of this work are as follows.

1. We utilize a general and systematic framework (column generation) to construct new asymmetric boosting algorithms, which can be applied to a variety of asymmetric losses. There is no heuristic strategy in our algorithms which may cause suboptimal solutions. In contrast, the global optimal solution is guaranteed for our algorithms.

Unlike Viola-Jones' asymmetric AdaBoost [3], the asymmetric effect of our methods spreads over the entire training process. The coefficients of all weak classifiers are updated at each iteration, which prevents the first weak classifier from absorbing the asymmetry. The number of weak classifiers does not need to be specified before training.

2. The asymmetric totally-corrective boosting algorithms introduce the asymmetric learning goal into both feature selection and ensemble classifier learning. Both the example weights and the linear classifier coefficients are learned in an asymmetric way.
3. In practice, L-BFGS-B [16] is used to solve the primal problem, which runs much faster than solving the dual problem and also less memory is needed.
4. We demonstrate that with the totally corrective optimization, the linear coefficients of some weak classifiers are set to zero by the algorithm such that fewer weak classifiers are needed. We present analysis on the theoretical condition and show how useful the historical information is for the training of successive nodes.

2 Asymmetric losses

In this section, we propose two asymmetric losses, which are motivated by asymmetric AdaBoost [3] and cost-sensitive LogitBoost [17], respectively.

We first introduce an asymmetric cost in the following form:

$$\mathrm{ACost} = \begin{cases} C_1 & \text{if } y = +1 \text{ and } \mathrm{sign}(F(\boldsymbol{x})) = -1, \\ C_2 & \text{if } y = -1 \text{ and } \mathrm{sign}(F(\boldsymbol{x})) = +1, \\ 0 & \text{if } y = \mathrm{sign}(F(\boldsymbol{x})). \end{cases}$$

Here $\boldsymbol{x}$ is the input data, y is the label and $F(\boldsymbol{x})$ is the learned classifier. Viola and Jones [3] directly take the product of ACost and the exponential loss $E_{\boldsymbol{X},Y}[\exp(-yF(\boldsymbol{x})]$ as the asymmetric loss:

$$E_{\boldsymbol{X},Y}[\big(\mathbf{I}(y=1)C_1 + \mathbf{I}(y=-1)C_2\big)\exp\big(-yF(\boldsymbol{x})\big)],$$

where $\mathbf{I}(\cdot)$ is the indicator function. In a similar manner, we can also form an asymmetric loss from the logistic loss $E_{\boldsymbol{X},Y}[\mathrm{logit}\big(-yF(\boldsymbol{x})\big)]$:

$$\mathrm{ALoss}_1 = E_{\boldsymbol{X},Y}[\big(\mathbf{I}(y=1)C_1 + \mathbf{I}(y=-1)C_2\big)\mathrm{logit}\big(yF(\boldsymbol{x})\big)], \tag{1}$$

where $\mathrm{logit}(x) = \log(1+\exp(-x))$ is the logistic loss function.

Masnadi-Shirazi and Vasconcelos [17] proposed cost-sensitive boosting algorithms which optimize different versions of cost-sensitive losses by the means of gradient descent. They proved that the optimal cost-sensitive predictor minimizes the expected loss:

$$-E_{\boldsymbol{X},Y}[\mathbf{I}(y=1)\log(\mathrm{p_c}(\boldsymbol{x})) + \mathbf{I}(y=-1)\log(1-\mathrm{p_c}(\boldsymbol{x}))],$$

$$\text{where} \quad \mathrm{p_c}(\boldsymbol{x}) = \frac{e^{\gamma F(x)+\eta}}{e^{\gamma F(x)+\eta} + e^{-\gamma F(x)-\eta}}, \text{with } \gamma = \frac{C_1+C_2}{2}, \eta = \frac{1}{2}\log\frac{C_2}{C_1}.$$

With fixing γ to 1, the expected loss can be reformulated to

$$\mathrm{ALoss}_2 = E_{\boldsymbol{X},Y}[\mathrm{logit}\big(yF(\boldsymbol{x}) + 2y\eta\big)]. \tag{2}$$

3 Asymmetric totally-corrective boosting

In this section, we construct asymmetric totally-corrective boosting algorithms (termed AsymBoost$_{\mathrm{TC}}$ here) from the losses (1) and (2) discussed previously. In contrast to the methods constructing boosting-like algorithms in [17], [18] and [19], we use column generation to design our totally corrective boosting algorithms, inspired by [20] and [5].

Suppose there are M training examples (M_1 positives and M_2 negatives), and the sequence of examples are arranged according to the labels (positives first). The pool $\mathcal{H}$ contains N available weak classifiers. The matrix $H \in \mathbb{Z}^{M\times N}$ contains binary outputs of weak classifiers in $\mathcal{H}$ for training examples, namely $H_{ij} = h_j(x_i)$. We are aiming to learn a linear combination $F_{\boldsymbol{w}}(\cdot) = \sum_{j=1}^{N} w_i h_j(\cdot)$. C_1 and C_2 are costs for misclassifying positives and negatives, respectively. We assign the asymmetric factor $k = C_2/C_1$ and restrict $\gamma = (C_1 + C_2)/2$ to 1, thus C_1 and C_2 are fixed for a given k.

The problems of the two AsymBoost$_{\mathrm{TC}}$ algorithms can be expressed as:

$$\min_{\boldsymbol{w}} \sum_{i=1}^{M} l_i \mathrm{logit}(z_i) + \theta \mathbf{1}^\top \boldsymbol{w} \quad \text{s.t. } \boldsymbol{w} \succcurlyeq \mathbf{0},\ z_i = y_i H_i \boldsymbol{w}, \tag{3}$$

where $\boldsymbol{l} = [C_1/M_1, \cdots, C_2/M_2, \cdots]^\top$, and

$$\min_{\boldsymbol{w}} \sum_{i=1}^{M} e_i \mathrm{logit}(z_i + 2y_i\eta) + \theta \mathbf{1}^\top \boldsymbol{w} \quad \text{s.t. } \boldsymbol{w} \succcurlyeq \mathbf{0},\ z_i = y_i H_i \boldsymbol{w}, \tag{4}$$

where $\boldsymbol{e} = [1/M_1, \cdots, 1/M_2, \cdots]^\top$. In both (3) and (4), z_i stands for the margin of the i-th training example. We refer (3) as AsymBoost$_{\mathrm{TC1}}$ and (4) as AsymBoost$_{\mathrm{TC2}}$. Note that here the optimization problems are ℓ_1-norm regularized. It is possible to use other format of regularization such as the ℓ_2-norm.

First we introduce a fact that the Fenchel conjugate [21] of the logistic loss function $\mathrm{logit}(x)$ is

$$\mathrm{logit}^*(u) = \begin{cases} (-u)\log(-u) + (1+u)\log(1+u), 0 \geq u \geq -1; \\ \infty, \qquad\qquad\qquad\qquad\quad \text{otherwise.} \end{cases}$$

Now we derive the Lagrange dual [21] of AsymBoost$_{\mathrm{TC1}}$. The Lagrangian of (3) is

$$L(\underbrace{\boldsymbol{w}, \boldsymbol{z}}_{\text{primal}}, \underbrace{\boldsymbol{\lambda}, \boldsymbol{u}}_{\text{dual}}) = \sum_{i=1}^{M} l_i \mathrm{logit}(z_i) + \theta \mathbf{1}^\top \boldsymbol{w} - \boldsymbol{\lambda}^\top \boldsymbol{w} + \sum_{i=1}^{M} u_i (z_i - y_i H_i \boldsymbol{w}).$$

The dual function

$$\begin{aligned} g(\boldsymbol{\lambda}, \boldsymbol{u}) &= \inf_{\boldsymbol{w},\boldsymbol{z}} L(\boldsymbol{w}, \boldsymbol{z}, \boldsymbol{\lambda}, \boldsymbol{u}) \\ &= -\sum_{i=1}^{M} \underbrace{\sup_{z_i} \Big(-u_i z_i - l_i \mathrm{logit}(z_i) \Big)}_{l_i \mathrm{logit}^*(-u_i/l_i)} + \underbrace{\inf_{\boldsymbol{w}} \Big(\theta \mathbf{1}^\top - \boldsymbol{\lambda}^\top - \sum_{i=1}^{M} u_i y_i H_i \Big)}_{\text{must be } \mathbf{0}} \boldsymbol{w}. \end{aligned}$$

The dual problem is

$$\begin{aligned} \max_{\boldsymbol{u}} \ & -\sum_{i=1}^{M}\Big[u_i\log(u_i)+(l_i-u_i)\log(l_i-u_i)\Big] \\ \text{s.t.} \ & \sum_{i=1}^{M} u_i y_i H_i \preccurlyeq \theta\mathbf{1}^\top,\ 0\preccurlyeq \boldsymbol{u}\preccurlyeq \boldsymbol{l}. \end{aligned} \tag{5}$$

Since the problem (3) is convex and the Slater's conditions are satisfied [21], the duality gap between the primal (3) and the dual (5) is zero. Therefore, the solutions of (3) and (5) are the same. Through the KKT condition, the gradient of Lagrangian (5) over primal variable $\boldsymbol{z}$ and dual variable $\boldsymbol{u}$ should vanish at the optimum. Therefore, we can obtain the relationship between the optimal value of $\boldsymbol{z}$ and $\boldsymbol{u}$:

$$u_i^* = \frac{l_i\exp(-z_i^*)}{1+\exp(-z_i^*)}. \tag{6}$$

Similarly, we can get the dual problem of AsymBoost$_{\mathrm{TC2}}$, which is expressed as:

$$\begin{aligned} \max_{\boldsymbol{u}} \ & -\sum_{i=1}^{M}\Big[u_i\log(u_i)+(e_i-u_i)\log(e_i-u_i)+2u_iy_i\eta\Big] \\ \text{s.t.} \ & \sum_{i=1}^{M} u_i y_i H_i \preccurlyeq \theta\mathbf{1}^\top,\ 0\preccurlyeq \boldsymbol{u}\preccurlyeq \boldsymbol{e}, \end{aligned} \tag{7}$$

with

$$u_i^* = \frac{e_i\exp(-z_i^*-2y_i\eta)}{1+\exp(-z_i^*-2y_i\eta)}. \tag{8}$$

In practice, the total number of weak classifiers, N, could be extremely large, so we can not solve the primal problems (3) and (4) directly. However equivalently, we can optimize the duals (5) and (7) iteratively using column generation [20]. In each round, we add the most violated constraint by finding a weak classifier satisfying:

$$h^\star(\cdot) = \operatorname*{argmax}_{h(\cdot)} \sum_{i=1}^{M} u_i y_i h(\boldsymbol{x}_i). \tag{9}$$

This step is the same as training a weak classifier in AdaBoost and LPBoost, in which one tries to find a weak classifier with the maximal edge (*i.e.* the minimal weighted error). The *edge* of h_j is defined as $\sum_{i=1}^{M} u_i y_i h_j(x_i)$, which is the inverse of the weighted error. Then we solve the restricted dual problem with one more constraint than the previous round, and update the linear coefficients of weak classifiers ($\boldsymbol{w}$) and the weights of training examples ($\boldsymbol{u}$). Adding one constraint into the dual problem corresponds to adding one variable into the primal problem. Since the primal problem and dual problem are equivalent, we can either solve the restricted dual or the restricted primal in

Algorithm 1 The training algorithms of AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$.

Input: A training set with M labeled examples (M_1 positives and M_2 negatives); termination tolerant $\varepsilon > 0$; regularization parameter θ; asymmetric factor k; maximum number of weak classifiers $N_{\max}$.

Initialization: $N = 0$; $\boldsymbol{w} = \mathbf{0}$; and $u_i = l_i/2$ or $e_i/(1 + k^{-y_i})$, $i = 1 \cdots M$.

for iteration $= 1 : N_{\max}$ **do**

- Train a weak classifier $h'(\cdot) = \mathrm{argmmax}_{h(\cdot)} \sum_{i=1}^{M} u_i y_i h(\boldsymbol{x}_i)$.
- Check for the termination condition:
  if iteration > 1 and $\sum_{i=1}^{M} u_i y_i h'(\mathbf{x}_i) < \theta + \varepsilon$, **then** break;
- Increment the number of weak classifiers $N = N + 1$.
- Add $h'(\cdot)$ to the restricted master problem;
- Solve the primal problem (3) or (4) (or the dual problem (5) or (7)) and update u_i ($i = 1 \cdots M$) and w_j ($j = 1 \cdots N$) .

Output: The selected weak classifiers are $h_1, h_2, \ldots, h_N$. The final strong classifier is: $F(\boldsymbol{x}) = \sum_{j=1}^{N} w_j h_j(\boldsymbol{x})$.

practice. The algorithms of AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ are summarized in Algorithm 1. Note that, in practice, in order to achieve specific false negative rate (FNR) or false positive rate (FPR), an offset b is needed to be added into the final strong classifier: $F(\boldsymbol{x}) = \sum_{j=1}^{n} w_i h_j(\mathbf{x}) - b$, which can be obtained by a simple line search. The new weak classifier $h'(\cdot)$ corresponds to an extra variable to the primal and an extra constraint to the dual. Thus, the minimal value of the primal decreases with growing variables, and the maximal value of the dual problem also decreases with growing constraints. Furthermore, as the optimization problems involved are convex, Algorithm 1 is guaranteed to converge to the global optimum.

Next we show how AsymBoost$_{\mathrm{TC}}$ introduces the asymmetric learning into feature selection and ensemble classifier learning. Decision stumps are the most commonly used type of weak classifiers, and each stump only uses one dimension of the features. So the process of training weak classifiers (decision stumps) is equivalent to feature selection. In our framework, the weak classifier with the maximum edge (*i.e*. the minimal weighted error) is selected. From (6) and (8), the weight of i-th example, namely u_i, is affected by two factors: the asymmetric factor k and the current margin z_i. If we set $k = 1$, the weighting strategy goes back to being symmetric. On the other hand, the coefficients of the linear classifier, namely $\boldsymbol{w}$, are updated by solving the restricted primal problem at each iteration. The asymmetric factor k in the primal is absorbed by all the weak classifiers currently learned. So feature selection and ensemble classifier learning both consider the asymmetric factor k.

The number of variables of the primal problem is the number of weak classifiers, while for the dual problem, it is the number of training examples. In the cascade classifiers for face detection, the number of weak classifiers is usually much smaller than the number of training examples, so solving the primal is much cheaper than solving the dual. Since the primal problem has only simple box-bounding constraints, we can employ L-BFGS-B [16] to solve it. L-BFGS-B is a tool based on the quasi-Newton method for bound-constrained optimization.

Instead of maintaining the Hessian matrix, L-BFGS-B only needs the recent several updates of values and gradients for the cost function to approximate the Hessian matrix. Thus, L-BFGS-B requires less memory when running. In column generation, we can use the results from previous iteration as the starting point of current problem, which leads to further reductions in computation time.

The complementary slackness condition [21] suggests that $\lambda_j w_j = 0$. So we can get the conditions of sparseness:

$$\text{If } \lambda = \theta - \sum_{i=1}^{M} u_i y_i H_{i,j} > 0, \text{then } w_j = 0. \tag{10}$$

This means that, if the weak classifier $h_j(\cdot)$ is so "weak" that its edge is less than θ under the current distribution $\boldsymbol{u}$, its contribution to the ensemble classifier is "zero". From another viewpoint, the ℓ_1-norm regularization term in the primal (3) and (4), leads to a sparse result. The parameter θ controls the degree of the sparseness. The larger θ is, the sparser the result would be.

4 Experiments

4.1 Results on synthetic data

To show the behavior of our algorithms, we construct a 2D data set, in which the positive data follow the 2D normal distribution ($N(0, 0.1\mathbf{I})$), and the negative data form a ring with uniformly distributed angles and normally distributed radius ($N(1.0, 0.2)$). Totally 2000 examples are generated (1000 positives and 1000 negatives), 50% of data for training and the other half for test. We compare AdaBoost, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ on this data set. All the training processes are stopped at 100 decision stumps. For AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$, we fix θ to 0.01, and use a group of k's $\{1.2, 1.4, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.8, 3.0\}$.

From Figures 1 (1) and (2), we find that the larger k is, the bigger the area for positive output becomes, which means that the asymmetric LogitBoost tends to make a positive decision for the region where positive and negative data are mixed together. Another observation is that AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ have almost the same decision boundaries on this data set with same k's.

Figures 1 (3) and (4) demonstrate trends of false rates with the growth of asymmetric factor (k). The results of AdaBoost is considered as the baseline. For all k's, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ achieve lower false negative rates and higher false positive rates than AdaBoost. With the growth of k, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ become more aggressive to reduce the false negative rate, with the sacrifice of a higher false positive rate.

4.2 Face detection

We collect 9832 mirrored frontal face images and about 10115 large background images. 5000 face images and 7000 background images are used for training, and 4832 face images and 3115 background images for validation. Five basic types of Haar features are calculated on each 24×24 image, and totally generate 162336 features. Decision stumps on those 162336 features construct the pool of weak classifiers.

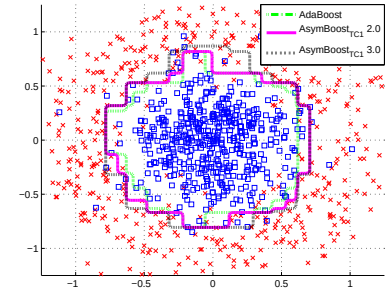

(1) AsymBoost$_{\mathrm{TC1}}$ vs AdaBoost

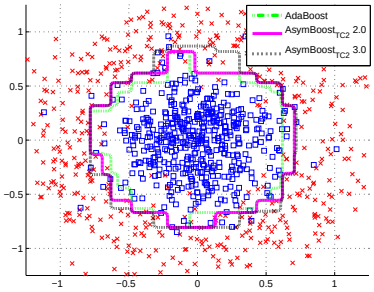

(2) AsymBoost$_{\mathrm{TC2}}$ vs AdaBoost

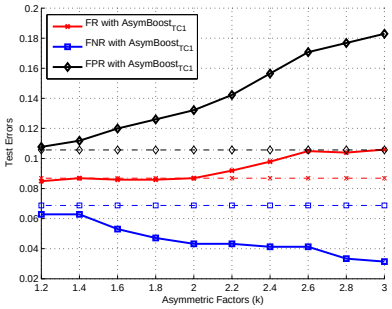

(3) False rates for AsymBoost$_{\mathrm{TC1}}$

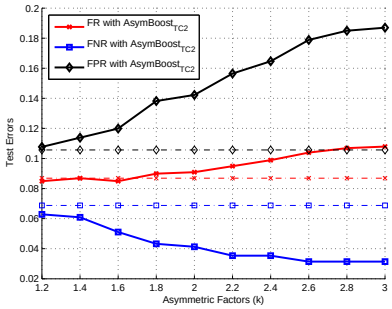

(4) False rates for AsymBoost$_{\mathrm{TC2}}$

Fig. 1: Results on the synthetic data for AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$, with a group of asymmetric factor ks. As the baseline, the results for AdaBoost are also shown in these figures. (1) and (2) demonstrate decision boundaries learned by AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$, with k is 2.0 or 3.0. The $\times$'s and $\square$'s stand for training negatives and training positives respectively. (3) and (4) demonstrate false rates (FR), false positive rates (FPR) and false negative rates (FNR) on test set with a group of ks (1.2, 1.4, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.8 or 3.0), and the corresponding rates for AdaBoost is shown as dashed lines.

Single-node detectors Single-node classifiers with AdaBoost, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ are trained. The parameters θ and k are simply set to 0.001 and 7.0. 5000 faces and 5000 non-faces are used for training, while 4832 faces and 5000 non-faces are used for test. The training/validation non-faces are randomly cropped from training/validation background images.

Figure 2 (1) shows curves of detection rate with the false positive rate fixed at 0.25, while curves of false positive rates with 0.995 detection rate are shown in Figure 2 (2). We set the false positive rate fixed to 0.25 rather than the commonly used 0.5 in order to slow down the increasing speed of detection rates, otherwise detection rates would converge to 1.0 immediately. The increasing/decreasing speed of detection rate/false positive rate is faster than reported in [8] and [9]. The reason is possibly that we use

10000 examples for training and 9832 for testing, which are smaller than the data used in [8] and [9] (18000 training examples and 15000 test examples). We can see that under both situations, our algorithms achieve better performances than AdaBoost in most cases.

The benefits of our algorithms can be expressed in two-fold: (1) Given the same learning goal, our algorithms tend to use smaller number of weak classifiers. For example, from Figure 2 (2), if we want a classifier with a 0.995 detection rate and a 0.2 false positive rate, AdaBoost needs at least 43 weak classifiers while AsymBoost$_{\mathrm{TC1}}$ needs 32 and AsymBoost$_{\mathrm{TC2}}$ needs only 22. (2) Using the same number of weak classifiers, our algorithms achieve a higher detection rate or a lower false positive rate. For example, from Figure 2 (2), using 30 weak classifiers, both AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ achieve higher detection rates (0.9965 and 0.9975) than AdaBoost (0.9945).

Complete detectors Secondly, we train complete face detectors with AdaBoost, asymmetric-AdaBoost, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$. All detectors are trained using the same training set. We use two types of cascade framework for the detector training: the traditional cascade of Viola and Jones [2] and the multi-exit cascade presented in [11]. The latter utilizes decision information of previous nodes when judging instances in the current node. For fair comparison, all detectors use 24 nodes and 3332 weak classifiers. For each node, 5000 faces + 5000 non-faces are used for training, and 4832 faces + 5000 non-faces are used for validation. All non-faces are cropped from background images. The asymmetric factor k for asymmetric-AdaBoost, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ are selected from $\{1.2, 1.5, 2.0, 3.0, 4.0, 5.0, 6.0\}$. The regularization factor θ for AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ are chosen from $\{\frac{1}{50}, \frac{1}{60}, \frac{1}{70}, \frac{1}{80}, \frac{1}{90}, \frac{1}{100}, \frac{1}{200}, \frac{1}{400}, \frac{1}{800}, \frac{1}{1000}\}$. It takes about four hours to train a AsymBoost$_{\mathrm{TC}}$ face detector on a machine with 8 Intel Xeon E5520 cores and 32GB memory. Comparing with AdaBoost, only around 0.5 hour extra time is spent on solving the primal problem at each iteration. We can say that, in the context of face detection, the training time of AsymBoost$_{\mathrm{TC}}$ is nearly the same as AdaBoost.

ROC curves on the CMU/MIT data set are shown in Figure 3. Those images containing ambiguous faces are removed and 120 images are retained. From the figure, we can see that, asymmetric-AdaBoost outperforms AdaBoost in both Viola-Jones cascade and multi-exit cascade, which coincide with what reported in [3]. Our algorithms have better performances than all other methods in all points and the improvements are more significant when the false positives are less than 100, which is the most commonly used region in practice.

As mentioned in the previous section, our algorithms produce sparse results to some extent. Some linear coefficients are zero when the corresponding weak classifiers satisfy the condition (10). In the multi-exit cascade, the sparse phenomenon becomes more clear. Since correctly classified negative data are discarded after each node is trained, the training data for each node are different. The "closer" nodes share more common training examples, while the nodes "far away" from each other have distinct training data. The greater the distance between two nodes, the more uncorrelated they become. Therefore, the weak classifiers in the early nodes may perform poorly on the last node, thus tending to obtain zero coefficients. We call those weak classifiers with non-zero

Table 1: The ratio of weak classifiers selected at the i-th node (column) appearing with non-zero coefficients in the j-th node (row). The ratios decrease along with the growth of the node index in each column.

Node Index	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	1.00														
2	1.00	1.00													
3	1.00	1.00	1.00												
4	0.86	1.00	0.97	1.00											
5	0.43	0.93	0.97	0.97	1.00										
6	0.71	0.93	0.90	1.00	0.96	1.00									
7	0.43	0.87	0.87	0.97	0.92	0.92	1.00								
8	0.29	0.40	0.70	0.73	0.74	0.88	0.74	1.00							
9	0.00	0.27	0.50	0.60	0.76	0.72	0.66	0.67	1.00						
10	0.14	0.27	0.43	0.60	0.62	0.70	0.62	0.66	0.60	1.00					
11	0.00	0.20	0.33	0.50	0.52	0.54	0.60	0.59	0.56	0.48	1.00				
12	0.14	0.20	0.40	0.40	0.56	0.50	0.54	0.61	0.55	0.46	0.36	1.00			
13	0.00	0.13	0.33	0.37	0.36	0.54	0.40	0.47	0.47	0.46	0.43	0.25	1.00		
14	0.00	0.07	0.17	0.40	0.28	0.50	0.42	0.49	0.50	0.53	0.45	0.43	0.35	1.00	
15	0.00	0.13	0.20	0.27	0.36	0.38	0.46	0.41	0.52	0.42	0.49	0.44	0.34	0.27	1.00

Table 2: Comparison of the numbers of the effective weak classifiers for the stage-wise boosting (SWB) and the totally-corrective boosting (TCB). We take AdaBoost and AsymBoost$_{\mathrm{TC1}}$ as representative types of SWB and TCB, both of which are trained in the multi-exit cascade for face detection.

Node Index	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
SWB	7	22	52	82	132	182	232	332	452	592	752	932	1132	1332	1532	1732	1932	2132
TCB	7	22	52	80	125	174	213	269	331	441	464	538	570	681	717	744	742	879

coefficients "effective" weak classifiers. Table 1 shows the ratios of "effective" weak classifiers contributed by one node to a specific successive node. To save space, only the first 15 nodes are demonstrated. We can see that, the ratio decreases with the growth of the node index, which means that the farther the preceding node is from the current node, the less useful it is for the current node. For example, the first node has almost no contribution after the eighth node. Table 2 shows the number of effective weak classifiers used by our algorithm and the traditional stage-wise boosting. All weak classifiers in stage-wise boosting have non-zero coefficients, while our totally-corrective algorithm uses much less effective weak classifiers.

5 Conclusion

We have proposed two asymmetric totally-corrective boosting algorithms for object detection, which are implemented by the column generation technique in convex opti-

mization. Our algorithms introduce asymmetry into both feature selection and ensemble classifier learning in a systematic way.

Both our algorithms achieve better results for face detection than AdaBoost and Viola-Jones' asymmetric AdaBoost. An observation is that we can not see great differences on performances between AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$ in our experiments. For the face detection task, AdaBoost already achieves a very promising result, so the improvements of our method are not very significant.

One drawback of our algorithms is there are two parameters to be tuned. For different nodes, the optimal parameters should not be the same. In this work, we have used the same parameters for all nodes. Nevertheless, since the probability of negative examples decreases with the node index, the degree of the asymmetry between positive and negative examples also deceases. The optimal k may decline with the node index.

The framework for constructing totally-corrective boosting algorithms is general, so we can consider other asymmetric losses (*e.g.*, asymmetric exponential loss) to form new asymmetric boosting algorithms. In column generation, there is no restriction that only one constraint is added at each iteration. Actually, we can add several violated constraints at each iteration, which means that we can produce multiple weak classifiers in one round. By doing this, we can speed up the learning process.

Motivated by the analysis of sparseness, we find that the very early nodes contribute little information for training the later nodes. Based on this, we can exclude some useless nodes when the node index grows, which will simplify the multi-exit structure and shorten the testing time.

References

1. Paisitkriangkrai, S., Shen, C., Zhang, J.: Fast pedestrian detection using a cascade of boosted covariance features. IEEE Trans. Circuits Syst. Video Technol. **18** (2008) 1140–1151
2. Viola, P., Jones, M.J.: Robust real-time face detection. Int. J. Comp. Vis. **57** (2004) 137–154
3. Viola, P., Jones, M.: Fast and robust classification using asymmetric AdaBoost and a detector cascade. In: Proc. Adv. Neural Inf. Process. Syst., MIT Press (2002) 1311–1318
4. Paisitkriangkrai, S., Shen, C., Zhang, J.: Efficiently training a better visual detector with sparse Eigenvectors. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., Miami, Florida, US (2009)
5. Shen, C., Li, H.: On the dual formulation of boosting algorithms. IEEE Trans. Pattern Anal. Mach. Intell. (2010) Online 25 Feb. 2010. IEEE computer Society Digital Library. `http://doi.ieeecomputersociety.org/10.1109/TPAMI.2010.47`.
6. Shen, C., Wang, P., Li, H.: LACBoost and FisherBoost: Optimally building cascade classifiers. In: Proc. Eur. Conf. Comp. Vis. Volume 2., Crete Island, Greece, Lecture Notes in Computer Science (LNCS) 6312, Springer-Verlag (2010) 608–621
7. Fan, W., Stolfo, S., Zhang, J., Chan, P.: Adacost: Misclassification cost-sensitive boosting. In: Proc. Int. Conf. Mach. Learn. (1999) 97–105
8. Li, S.Z., Zhang, Z.: FloatBoost learning and statistical face detection. IEEE Trans. Pattern Anal. Mach. Intell. **26** (2004) 1112–1123
9. Xiao, R., Zhu, L., Zhang, H.: Boosting chain learning for object detection. In: Proc. IEEE Int. Conf. Comp. Vis. (2003) 709–715
10. Hou, X., Liu, C., Tan, T.: Learning boosted asymmetric classifiers for object detection. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2006)

11. Pham, M.T., Hoang, V.D.D., Cham, T.J.: Detection with multi-exit asymmetric boosting. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2008)
12. Bourdev, L., Brandt, J.: Robust object detection via soft cascade. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., San Diego, CA, US (2005) 236–243
13. Xiao, R., Zhu, H., Sun, H., Tang, X.: Dynamic cascades for face detection. In: Proc. IEEE Int. Conf. Comp. Vis., Rio de Janeiro, Brazil (2007)
14. Wu, J., Brubaker, S.C., Mullin, M.D., Rehg, J.M.: Fast asymmetric learning for cascade face detection. IEEE Trans. Pattern Anal. Mach. Intell. **30** (2008) 369–382
15. Liu, C., Shum, H.Y.: Kullback-Leibler boosting. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. Volume 1., Madison, Wisconsin (2003) 587–594
16. Zhu, C., Byrd, R.H., Nocedal, J.: L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. ACM Trans. Mathematical Software **23** (1997) 550–560
17. Masnadi-Shirazi, H., Vasconcelos, N.: Cost-sensitive boosting. IEEE Trans. Pattern Anal. Mach. Intell. (2010)
18. Friedman, J., Hastie, T., Tibshirani, R.: Additive logistic regression: a statistical view of boosting. Ann. Statist. **28** (2000) 337–407
19. Rätsch, G., Mika, S., Schölkopf, B., Müller, K.R.: Constructing boosting algorithms from SVMs: An application to one-class classification. IEEE Trans. Pattern Anal. Mach. Intell. **24** (2002) 1184–1199
20. Demiriz, A., Bennett, K., Shawe-Taylor, J.: Linear programming boosting via column generation. Mach. Learn. **46** (2002) 225–254
21. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press (2004)

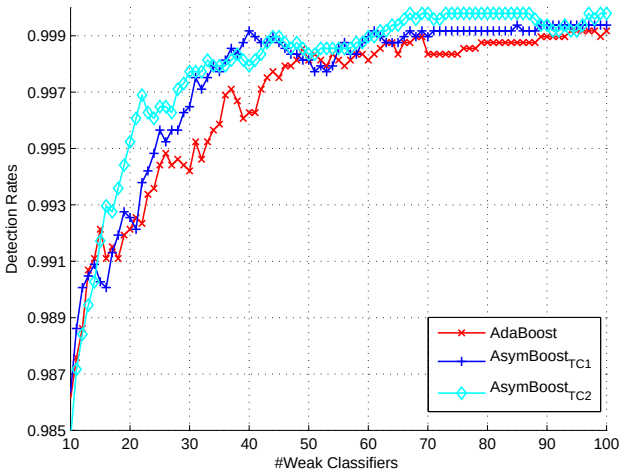

(1) DR with fixed FPR

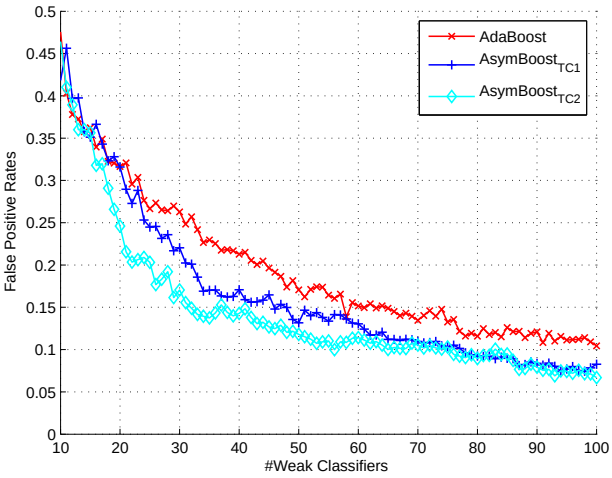

(2) FPR with fixed DR

Fig. 2: Testing curves of single-node classifiers for AdaBoost, AsymBoost$_{\mathrm{TC1}}$ and AsymBoost$_{\mathrm{TC2}}$. All the classifiers use the same training and test data sets. (1) shows curves of detection rates (DR) with false positive rates (FPR) fixed to 0.25, (2) shows curves of FPR with DR fixed to 0.995. FPR or DR are evaluated at each weak classifier.

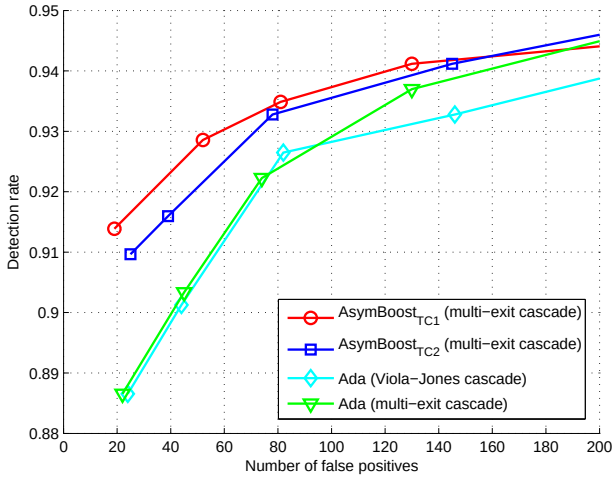

(1)

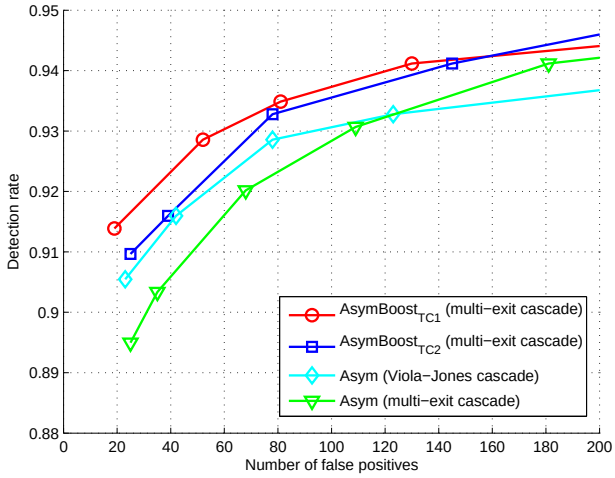

(2)

Fig. 3: Performances of cascades evaluated by ROC curves on the MIT+CMU data set. AdaBoost is referred to “Ada”, and Asymmetric AdaBoost [2] is referred to “Asym”. “Viola-Jones cascade” means the traditional cascade used in [3] .